

Distributed mining of large scale remote sensing image archives on public computing infrastructures

Luigi Mascolo, Marco Quartulli, Pietro Guccione, Giovanni Nico and Igor G. Olaizola

Abstract

Earth Observation (EO) mining aims at supporting efficient access to and exploration of petabyte-scale space- and airborne remote sensing archives that are currently expanding at rates of terabytes per day. A significant challenge is performing the analysis required by envisaged applications — like for instance process mapping for environmental risk management — in reasonable time. In this work, we address the problem of content-based image retrieval via example-based queries from EO data archives. In particular, we focus on the analysis of polarimetric SAR data, for which target decomposition theorems have proved fundamental in discovering patterns in data and characterize the ground scattering properties. To this end, we propose an interactive region-oriented content-based image mining system in which 1) unsupervised ingestion processes are distributed onto virtual machines in elastic, on-demand computing infrastructures 2) archive-scale content hierarchical indexing is implemented in terms of a “big data” analytics cluster-computing framework 3) query processing amounts to traversing the generated binary tree index, computing distances that correspond to descriptor-based similarity measures between image groups and a query image tile. We describe in depth both the strategies and the actual implementations for the ingestion and indexing components, and verify the approach by experiments carried out on the NASA/JPL UAVSAR full polarimetric data archive.

We report the results of the tests performed on computer clusters by using a public Infrastructure-as-a-Service and evaluating the impact of cluster configuration on system performance. Results are promising for data mapping and information retrieval applications.

Keywords

Content-Based Retrieval, Remote Sensing, Cloude-Pottier decomposition, Cloud Computing, Big Data.

I. INTRODUCTION

REMOTELY sensed data volumes are growing at faster and faster rates due to the increasing number of spaceborne and airborne Earth Observation (EO) missions and to the improvement of image resolution. As an example, according to a yearly report published at the end of 2013, the archives of NASA’s Earth Observing System Data and Information System (EOSDIS) have a volume of almost 10 petabytes, with 6,900 accessible datasets and an average archive growth of approximately 8.5 terabytes per day. The possibility of accessing and processing large databases of remote sensing images allows planetary scale applications like deforestation monitoring, glacial retreat investigation, urban development mapping, land cover classification and so on. Allowing an efficient discovery, annotation and retrieval of data products is the goal of EO mining systems [13].

Common methods to search for images are based on meta-data, that is a user formulates a query in terms of geographic location, sensor parameters, time of acquisition, manual annotations, and the system returns images that are relevant to the query. The problem with such approaches is that they hardly satisfy user needs in terms of human perception, while manually annotating large volumes of images for describing their content is so expensive and time-consuming that it becomes impossible when facing real scenarios, where petabyte-scale image archives are the rule. An alternative approach is Content-Based Image Retrieval (CBIR), where a description of an image in terms of primitive features (textures, shapes, relevant colors etc...) is automatically extracted from its content, providing the user with query and retrieval methods that are closer to visual perception. Despite several systems for content-based retrieval having been proposed for mining EO image databases, most of them have not been conceived for scaling to massive volumes of data and it is estimated that up to 95% of the data present in existing archives has never been accessed or analyzed by a human analyst [9].

The fundamental limitation of remote sensing content information retrieval has been about their applicability to extended archives: operational scalability has not usually been a characteristic of a number of prototype systems presented in the literature. A system for automatic object and tile-based feature extraction, content mining and indexing has been proposed in [16], [8]. Experiments are conducted on a 40 Gigabyte archive of high resolution images from the IKONOS and QuickBird satellites. Multiple machines are used to answer user queries, with each machine managing a different feature related index and processing queries in parallel. When queried, the system answers by combining the results from the distributed index. However, in substance, both the system architecture and the algorithms are single node-based, which is unlikely to provide the computational capabilities required to process millions or billions of records. Ingestion and, in particular, indexing procedures

L. Mascolo and P. Guccione are with Department of Electrical and Information Engineering, Polytechnic of Bari, Via E. Orabona 4, 70125 Bari, Italy e-mail: lui.mascolo@gmail.com.

M. Quartulli and I. G. Olaizola are with Vicomtech-IK4, Mikeletegi Pasealekua 57, 20009 Donostia-San Sebastián, Spain e-mail: mquartulli@vicomtech.org. G. Nico is with Istituto per le Applicazioni del Calcolo, National Research Council, Via Amendola 122, 70126 Bari, Italy.

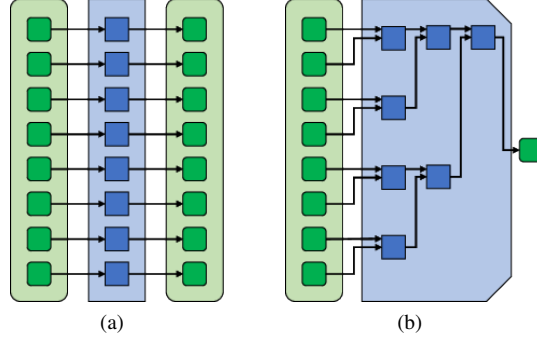


Fig. 1. Concept diagrams: parallel implementations of a map pattern and a reduce pattern. In the map pattern shown in Fig 1a, an elemental function is independently applied over every element of a data collection. The reduce pattern performs pairwise combinations of the elements of the collection in order to create a summary value, as shown in Fig 1b.

are considered offline operations. In [12] a similarity-based retrieval system for Terabyte scale image archives is presented. The reported experiments are conducted on a large image dataset whose descriptors are stored on a cluster of commodity machines. However parallel processing procedures needed for the generation of the descriptors are not considered, and index creation is performed by building a tree-structured index on a single machine, relying on the descriptors of a limited subset of archived products obtained by randomly sampling the entire collection. Distributing index creation procedures across multiple machines is in fact a challenging task, consisting of long-running data- and computational-intensive operations for which scalable data analysis algorithms are needed.

In this paper, we propose an EO data search framework that allows the exploration of massive datasets via example-based queries. We show that the system can be exploited to interactively query large-scale EO imagery archives with logarithmic search complexity and report results of tests performed by using a public Infrastructure-as-a-Service. Without imposing any constraint on the extension of the system to any kind of sensor, we focus in this specific work on polarimetric Synthetic Aperture Radar (SAR) data, where the improvements in the sensor resolution and the use of multiple polarizations are causing a relevant growth of data volume. The publicly available database of polarimetric SAR images collected by the UAVSAR [14] sensor has been considered as a case study.

The paper is structured as follows. In Section II, the analysis and parallelization methodologies involved in different parts of the system architecture are introduced. In Section III, the basic concepts of target decomposition theory for polarimetric SAR data are introduced and the specific technique employed in the system, Cloude-Pottier decomposition, is briefly described. In Section IV the system architecture and the procedures adopted for the data ingestion and the index creation are outlined. The results of the experiments carried out by using cloud-based computing infrastructures are reported in Section V. Finally, in Section VI, a discussion on the results and the conclusions close the article.

II. ANALYSIS AND PARALLELIZATION METHODOLOGY

CBIR systems are typically based on three separate data processing phases and on a separate client UI.

The first phase entails an unsupervised ingestion step, in which single archive products are analyzed independently of each other in order to characterize them based on their contents. The second phase consists of an indexing step in which the obtained characterizations are compared with each other in order to organize the contents of the archive in groups of similar characteristics, in order to simplify navigation. This phase is unsupervised as well, and results in the production of a content index. A last and interactive phase is query processing, in which the unsupervised characterization of an example query is used to navigate the index tree to retrieve the most relevant archive items.

In general terms, the ingestion phase can be seen as “perfectly parallel” or “embarrassingly parallel” [5] and fits a single instruction, multiple data paradigm, while the second phase, index formation, requires iterative processes in which archive elements are repeatedly compared with each other. The third phase, query processing, already operates on efficient data structures and often does not require specific parallelization methodologies.

Perfectly parallel workloads typically involve multiple tasks that have no dependency on one another. It is in this sense that the first phase can be expressed by a *map* pattern (see Fig. 1a) that invokes an elemental function over every element of a data collection in parallel [11]. A fundamental requirement for the implemented elemental functions is to not modify global data on which other instances of that function depend. Examples of operations amenable to direct implementation as map-like functions that are typically part of the ingestion phase of an EO image mining system include data product access, tiling and feature extraction. All such functions operate on isolated elements of the input archive, but their usage is generally not uniform in time. Virtual resources in elastic cloud computing infrastructures are well suited to such a scenario, allowing to distribute the ingestion processes dynamically and on-demand, achieving economies of scale and maximizing the effectiveness of the shared resources.

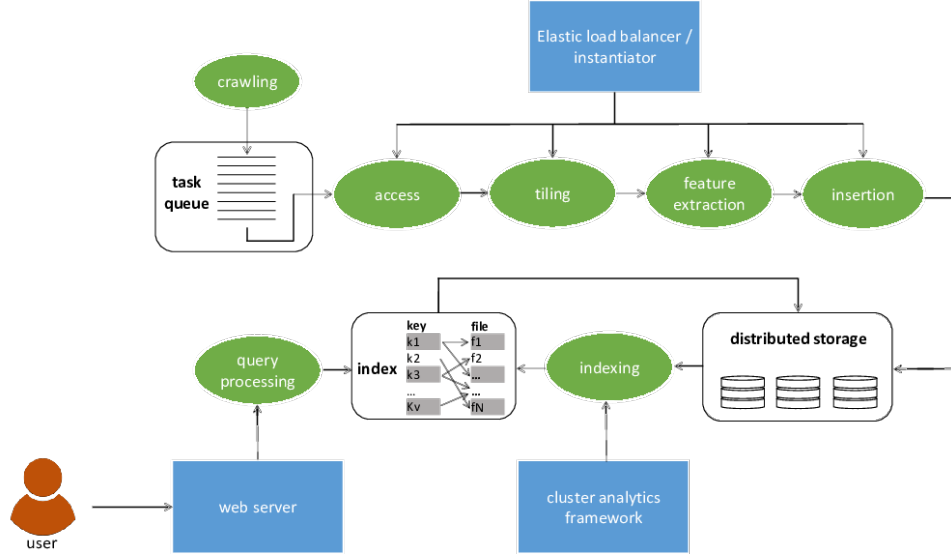


Fig. 2. A sketch of the high level architectural decomposition of the system. In the ingestion phase, in the upper part of the diagram, crawlers populate a task queue with identifiers of products to be ingested. One or more workers instantiated by an elastic load balancing mechanism access the data products, divide them into tiles, extract descriptors for each of them (feature extraction) and insert these descriptors into a distributed storage system. In the indexing phase, in the bottom half of the diagram, a cluster of machines is orchestrated by a cluster analytics framework to build a content-based index of the items in the distributed storage database. The results of the indexing process consist of a data structure that allows logarithmic-complexity content queries to take place. They are made available to a query processing system that is invoked via a web server based on interface events generated by a user in an interactive visualization subsystem.

The scenario of the second phase is different: the dataset has to be seen as a whole rather than a collection of separate elements, as the iterative machine-learning processes involved require communication between tasks, especially when combining intermediate results. The combine function defines a *reduce* pattern, with the elements of a collection being combined in any order to create a summary (see Fig. 1b).

MapReduce is a programming model that is well suited to this scenario. It requires the definition of two user-defined methods that process data organized as key/value pairs. The map function takes in input key/value pairs and, for each of them, applies the mapping procedures, resulting in zero, one or multiple key/value pairs in output. The reduce function takes the intermediate key/value groups, i.e. the grouped-by-key/value pairs, and processes together the intermediate values associated with the same key. A key advantage of such programming model is the possibility of keeping the computation close to the data, avoiding the need of moving the data from the place where they are stored. Expressing an algorithm by sequences of map and reduce patterns allows implementing parallel applications in a way that is suitable for the execution on large clusters of commodity machines [4].

Based on these considerations and to demonstrate the feasibility of large scale remote sensing CBIR on inexpensive on-demand processing infrastructures, we have developed a complete system that allows to distribute both ingestion and indexing processes over computer clusters and is able to scale to large scale databases.

III. MINING POLARIMETRIC SAR DATA ARCHIVES

The potential of SAR polarimetry to characterize the physical properties of the Earth surface has led to a variety of applications that aim at exploiting the ground scattering mechanism to extract geophysical parameters and perform landcover classification. Such parameter signatures are known to allow ground target classification, and therefore represent a direct possibility for the development of a system to mine large scale remote sensing archives.

Electromagnetic radiation incident on a target with a given polarization (horizontal or vertical) produces a backscattered wave that has, in general, both horizontal and vertical polarizations. The way a target changes the backscattered wave polarization is described by the scattering matrix, which describes the transformation of the electric field between the incident and backscattered wave:

$$\begin{bmatrix} E_H^b \\ E_V^b \end{bmatrix} = \begin{bmatrix} S_{HH} & S_{HV} \\ S_{VH} & S_{VV} \end{bmatrix} \begin{bmatrix} E_H^i \\ E_V^i \end{bmatrix}, \quad (1)$$

where $\mathbf{E}^i = [E_H^i \ E_V^i]^T$ is the electric field of the incident wave, $\mathbf{E}^b = [E_H^b \ E_V^b]^T$ is the electric field of the backscattered wave and S_{HH} , S_{HV} , S_{VH} and S_{VV} are the four complex coefficients of the scattering matrix, which are variable and depend on the nature of the target [10].

The elements on the main diagonal of the scattering matrix, S_{HH} and S_{VV} , produce the power return in the copolarized channels and the elements S_{HV} and S_{VH} produce the power return in the cross-polarized channels.

There are several ways to represent the scattering properties of a target. The coherence matrix is a possible power representation on which Cloude-Pottier decomposition theory relies [2]. Let us consider the vectorized version of the scattering matrix, based on Pauli spin elements:

$$\mathbf{k}_P = \frac{1}{\sqrt{2}} \begin{bmatrix} S_{HH} + S_{VV} \\ S_{HH} - S_{VV} \\ 2S_{HV} \end{bmatrix}, \quad (2)$$

where the reciprocity assumption $S_{HV} = S_{VH}$ has been supposed. The coherence matrix is a positive semi-definite Hermitian matrix, given by the product of the Pauli vector by itself:

$$\mathbf{T} = \mathbf{k}_P \mathbf{k}_P^T. \quad (3)$$

Cloude-Pottier decomposition relies on the eigen-decomposition of the (averaged) coherence matrix:

$$\mathbf{T} = \mathbf{U} \mathbf{\Sigma} \mathbf{U}^{-1} = \sum_{i=1}^3 \lambda_i \mathbf{u}_i \cdot \mathbf{u}_i^{T*}, \quad (4)$$

where $\mathbf{\Sigma} = \text{diag}(\lambda_1, \lambda_2, \lambda_3)$ is a diagonal matrix with non-negative real components, corresponding to the eigenvalues of \mathbf{T} , and \mathbf{U} is a square matrix whose columns are the eigenvectors of \mathbf{T} . Let us further account for the following parameterization of the eigenvectors, holding for the case of scattering medium that does not have azimuth symmetry:

$$\mathbf{u}_i = e^{j\phi_i} \begin{bmatrix} \cos \alpha_i \\ \sin \alpha_i \cos \beta_i e^{j\delta_i} \\ \sin \alpha_i \sin \beta_i e^{j\gamma_i} \end{bmatrix}. \quad (5)$$

Based on statistical considerations [10], the mean parameters of the dominant scattering mechanism can be extracted from the 3-by-3 coherency matrix as a mean unit target vector given by:

$$\mathbf{u}_0 = e^{j\phi} \begin{bmatrix} \cos \bar{\alpha} \\ \sin \bar{\alpha} \cos \bar{\beta} e^{j\bar{\delta}} \\ \sin \bar{\alpha} \sin \bar{\beta} e^{j\bar{\gamma}} \end{bmatrix}. \quad (6)$$

It has been shown that the mean parameter $\bar{\alpha}$, ranging from 0 to $\pi/2$ rad, is roll-invariant and that it is the main parameter for identifying the dominant ground scattering mechanism, while $\bar{\beta}$, $\bar{\delta}$ and $\bar{\gamma}$ can be used to define the target polarization orientation angle and ϕ is physically equivalent to an absolute target phase. The best estimate of $\bar{\alpha}$ is:

$$\bar{\alpha} = \sum_{k=1}^3 P_k \alpha_k, \quad (7)$$

with P_k representing pseudo probabilities:

$$P_k = \frac{\lambda_k}{\sum_{i=1}^3 \lambda_i}, \quad \sum_{k=1}^3 P_k = 1. \quad (8)$$

Based on the value of $\bar{\alpha}$, a characterization of the scatterer physical properties can be inferred, in terms of single (surface), volume or double bounce scattering, respectively corresponding to low, average and high values of $\bar{\alpha}$.

The coherence matrix eigen-decomposition can also be used to estimate polarimetric entropy H , a measure of the degree of statistical disorder (randomness) of the scatterer:

$$H = - \sum_{k=1}^N P_k \log_N (P_k), \quad (9)$$

with $N = 3$ for a monostatic radar and $N = 4$ for a bistatic radar. Polarimetric entropy ranges from 0 to 1. Low values of this parameter indicate weakly depolarizing targets (point scatterers), while high values characterize depolarizing targets (mixture of point scatterers). In the limit case of $H = 1$, the target scattering corresponds to a random noise process.

As a complementary information to the polarimetric entropy, the polarimetric anisotropy A can be estimated from the eigenvalues of the coherence matrix:

$$A = \frac{\lambda_2 - \lambda_3}{\lambda_2 + \lambda_3}. \quad (10)$$

Polarimetric anisotropy ranges from 0 to 1 and provides a measure of the relative importance of the second and the third eigenvalues of the eigen-decomposition. Both H and A are roll-invariant parameters, as the eigenvalues are rotationally invariant.

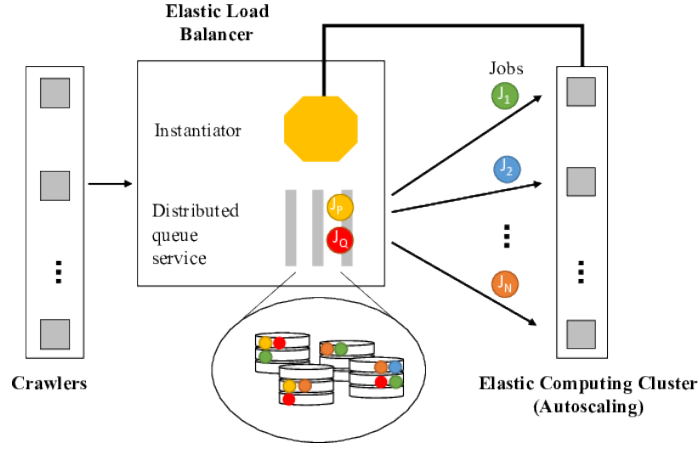


Fig. 3. A schema of the elastic load balancing system. Crawling agents insert pointers to new products in a distributed queue service to mark them for ingestion; jobs are made available to machines in the elastic computing cloud, whose resources are managed by an instantiator module.

IV. SYSTEM ARCHITECTURE

Four subsystems compose the architecture of the proposed CBIR system: data ingestion, indexing, query processing and interactive visualization. The system allows users to query large-scale archives of polarimetric data and retrieve image tiles with scattering characteristics similar to the ones of a user-provided query one. A schematic representation of the system is reported in Fig. 2.

The visualization subsystem provides the interactive environment by which the user — typically an image analyst — can communicate with the retrieval system and submit a query image tile. The query is unsupervisedly analyzed to extract one or more vectors of content-based features. The resulting query signature feature vectors are sent to the query execution subsystem, where the images with signatures most similar to the query image are efficiently retrieved. Finally, the retrieval results are shown to the user again by the user interface module.

The current section provides a detailed description of the architecture of the proposed CBIR system for polarimetric SAR data products, and analyzes the inner workings of each of the modules.

A. Data ingestion

Data ingestion is the first back-end module of the retrieval system. It consists of a system that retrieves images from the web or from a local database, extracts information about their contents in an unsupervised manner and stores it for later indexing.

Workloads in this stage are typically characterized by strong variability. For example, workload peaks are expected in early ingestion phases of a new database containing a large number of products. On the other hand, low demanding or idle states are also expected to happen in a full-time operating system. In order to avoid fast resource saturation or under utilization, the system has to be scalable, i.e. it has to be able to sustain both increasing and decreasing workloads with adequate performance by adapting hardware resources [7].

To this end, cloud-based computing environments can be exploited for scraping and performing processing and feature extraction operations on data products. In particular, elastic versions of such services can be implemented by allowing automated load balancers to on-demand provision and deprovision resources in the cloud environment. A load balancer allows the computing capacities to adapt to current requirements, i.e. to the amount of data that has to be ingested and processed at a certain time. Its operation is based on a distributed queue system and on a machine instance manager.

A further concrete issue are the heavy file transfers involved in this step in cases in which the storage and data analysis systems are far apart (as is of course the case of our prototype). This can easily bring to a fast saturation of networking resources, adding to the cost of ingestion.

1) *Crawling and elastic load balancing*: Three components are involved in the first phase of the ingestion system: the crawling agents, the elastic load balancer, which is in turn formed by a distributed queue system and a machine instance manager, and the elastic computing cluster. Crawling agents are charged with performing a continuous polling of the contents of the input data archive. In the specific case of the prototype under discussion, they perform an analysis of the NASA/JPL UAVSAR web pages with the aim of discovering new data products to ingest. When a new product is found, a pointer to it is created, and inserted in a distributed queue system. The distributed queue system is a scalable service that allows reliable queuing of independent jobs that cluster components have to perform. It acts as a buffer and communication channel among the crawling agents and the elastic cluster platform components, preventing the system to loose job instantiation requests when there is an excessive processing load or in case workers are scheduled to operate intermittently. The distributed nature of the queue also allows multiple interactions between crawlers and workers. In addition, a specific function of the queue manager

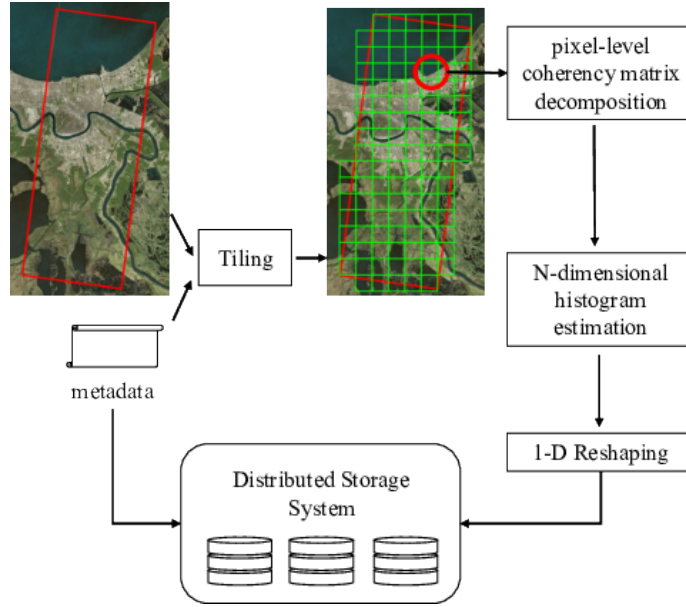


Fig. 4. Block scheme of tiling and feature extraction. The specific feature extraction procedure shown refers to scattering matrix decompositions applied to polarimetric data products.

is to allow better failure recovery strategies to be implemented in the system: when a worker node picks up a task to be executed, it also declares a timeout date after which it can be assumed that the processing will surely have been completed. If after this date the job has not been actively declared as completed by the worker tasked with executing it, the queue manager assumes that the processing failed and toggles a visibility flag on the scheduled job, which reappears in the queue for execution. The other component of the load balancing system is the worker machine instance manager, which is in charge of managing the cluster capacities in terms of both the number and the kind of instantiated machines based on explicitly defined scaling policies. Scaling policies can be based on simple metrics like the queue size, feedbacks about resource usage from subsequent modules in the ingestion chain or on more complex strategies for costs or energy saving that could entail intermittent service availability.

2) *Tiling and feature extraction*: Remote sensing image products cover vast areas of the Earth surface. Their content is therefore typically quite heterogeneous. Extracting a global description of each image product is not useful for most applications of content-based image retrieval systems, as, in general, the aim is to allow a user to identify specific targets or regions with similar characteristics.

The processing operations required for region-oriented content-based image mining consist of two main tasks: tiling the images in geo-referenced patches, whose size depends on the kind of sensor, on the resolution and on the aims of the content-based searches, and performing feature extraction operations.

The dimension of the tiles has to be chosen in such a way as to satisfy three main requirements:

- Preserving as much as possible the local characteristics of each patch of the original image product: a tile should be small enough to make it probable that it will include only one target class.
- Retaining enough information for the extraction of the ground target category. This general requirement means, in the context of polarimetric SAR images, retaining a number of pixels that is reasonable for an accurate estimation of the $H/\alpha/A$ density distribution;
- Allowing efficient use of the addressing space. Because memory devices are arrays of bytes, for the sake of memory allocation efficiency the tile dimension should be chosen following a power-of-2 scheme.

Furthermore, an overlap between the tiles can also be accounted for. During tile creations, image patches are univocally identified, georeferenced and processed for feature extraction. The block scheme of the operations implemented for polarimetric image ingestion is shown in Fig. 4.

3) *Distributed storage*: The operations in the first phase are typically disk I/O bound and need to operate on data products that have volumes of Gigabytes (around 10 Gigabytes of average size for UAVSAR ground-projected products). The obtained results need to be accessed globally for index formation. This naturally raises the problem of allowing distributed access to files in clustering environments. For this reason, the extracted image tile descriptors and metadata are pushed to a distributed file system.

A distributed file system allows storing information on networks of worker machines, offering at the same time scalability, fault tolerance and efficient computation performances. The distributed nature of the systems allows large datasets to be

Algorithm 1 Index generation

```
1: function INDEX(data)
2:    $model() \leftarrow cluster(data)$                                 # Initialize model (root node)
3:    $data \leftarrow model(data)$                                     # Label the dataset
4:    $leaves \leftarrow find\_leaves(model())$                         # Get leaves list
5:   while  $|leaves| \neq 0$  do                                       # Iterate for each leaf
6:      $leaf_{curr} \leftarrow take(leaves)$ 
7:     if not  $stopcriteria$  then
8:        $model() \leftarrow cluster(data, model(), leaf_{curr})$ 
9:        $data \leftarrow model(data)$ 
10:       $leaves \leftarrow find\_leaves(model())$ 
11:     else
12:        $leaves \leftarrow del\_leaf(leaves, leaf_{curr})$           # Prevent further bloomings
```

Fig. 5. Pseudo code of the tree growing procedure used by the indexing subsystem.




Fig. 6. Encoding/decoding scheme in Vector Quantization. The encoder maps input vectors to a finite set of codewords and sends the codeword index through the channel. The decoder holds a lookup table and is able to find the original codeword based on the received index.

accommodated on machines in the cluster. Data descriptors are stored in files that are split into one or more equal and fixed size chunks that are replicated and stored across the network nodes. Such operations allow the system to handle node failures, thereby preserving data from being lost and preventing the completion of long running computations on the cluster from depending on each single machine. Furthermore, distributed storage systems are designed to minimize the risk of network congestion and to increase overall system throughput. As network resources have bandwidth limits, data movements have to be discouraged, preferring a paradigm that allows bringing computation close to data and not vice-versa. Distributed file systems naturally fit such computational paradigms, as nodes working as storage units can also be employed as computational units in the cluster. We adopted for our system the Hadoop Distributed File System (HDFS) [15].

B. Retrieval and indexing

A crucial component for a large-scale CBIR system is an efficient methodology for the generation of a global index. In the proposed tile-based retrieval system, the index is the result of an unsupervised hierarchical clustering subsystem that has the purpose to provide an efficient way to group image tiles that are characterized by similar scattering behavior and that are therefore all likely to be relevant to a specific query. Without an index, the search engine would need to perform an $O(n)$ scan of the entire archive with unsustainable computational costs for mining even moderate size archives.

In a “big data” context, it is fundamental for the system to be able to perform indexing on large-scale datasets containing billions of entries, each representing the image content and eventually consisting of a large number of descriptors, and to allow fast and accurate retrieval results. In order to generate the content-based index, the system has to analyze the set of tile descriptors generated by the feature extraction module. Such analysis typically requires the application of machine learning algorithms. However, indexing large scale archives of data requires suitable implementations, able to distribute the computation among multiple processors in the cluster. To this purpose, we propose a mechanism for scalable indexing based on Tree-Structured Vector Quantization [6].

1) *Tree Structured Vector Quantization*: Vector Quantization is a signal processing technique that is popular in data compression. A generic schema of a vector quantizer consists of an encoding/decoding system. Let us represent the generic entry (a set of tile descriptors, in our case) in the database \mathbf{D} by a d -dimensional feature vector $\mathbf{x} = [x_1, x_2, \dots, x_d]$. The

encoder maps the d -dimensional vectors in the vector space \mathbb{R}^d to a finite set of vectors y_1, y_2, \dots, y_k , called code vectors or codewords, operating according to a nearest neighbor or minimum distortion rule. We here consider the squared error distortion, i.e. the square of the Euclidean distance between the input vector and the codeword:

$$d(\mathbf{x}, \mathbf{y}_i) = \|\mathbf{x} - \mathbf{y}_i\|_2^2 = \sum_{j=1}^d (x_j - y_{i,j})^2. \quad (11)$$

In this way, each codeword y_i has an associated nearest neighbor region, also called a Voronoi region, defined by:

$$V_i = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x} - \mathbf{y}_i\| \leq \|\mathbf{x} - \mathbf{y}_j\|, \forall j \neq i\}, \quad (12)$$

such that

$$\bigcup_{i=1}^k V_i = \mathbb{R}^d, \quad \forall i \neq j \quad (13)$$

and

$$\bigcap_{i=1}^k V_i = \emptyset, \quad \forall i \neq j. \quad (14)$$

Once a codeword is associated to an input d -dimensional vector, the corresponding codeword index is sent to the decoding system through a channel (a file system or a communication link, depending on the application). The decoder consists of a lookup table, i.e. a codebook containing all the possible codewords. When the index of a codeword is received, it will return the codeword corresponding to that index. The schema of the encoding/decoding system is shown in Fig. 6.

Several different approaches can be considered in the design a Vector Quantizer. Tree-Structured Vector Quantization (TSVQ) is a class of constrained structure quantizers [3]. In TSVQ, the codebook is constrained to have a tree structure. The encoder builds the codeword associated to an input vector by performing a sequence of binary comparisons, following a minimum distortion rule at each branch, until a leaf (terminal) node is reached. The path followed to reach the leaf node starting from the root indicates the binary sequence associated with the codeword.

A pseudo-code of the implementation of the tree-growing procedure is reported in IV-B. The tree structure is built starting from the root node. The root node is represented by the centroid of the entire set of feature vectors. From the root node, two new child nodes are estimated by applying the k -means clustering algorithm, each corresponding to the centroid of the space partitions $\mathbf{D} = \{D_1, D_2\}$ minimizing the within-cluster sum of squares cost function:

$$J_{WCSS} = \sum_{i=1}^2 \sum_{x \in D_i} \|\mathbf{x} - \mathbf{y}_i\|_2^2, \quad (15)$$

where \mathbf{y}_i represents the i th partition centroid. Then, each data point is assigned to its respective centroid by a binary string labeling, so that each group of data points defines a sub-tree. The process continues iterating on the discovered nodes, until a predefined stopping criterion is reached. The centroids of the items belonging to the leaves of the final tree structure, each with an associated binary string, represent the entries of the index. During the growing process, each split produces a decrease of the average distortion and an increase in the average length of the binary string. Based on these observations, several conditions are possible as stopping criteria. For example, the algorithm can stop growing a sub-tree (thereby defining a leaf node) when one of the following events occurs:

- the node contains a predefined minimum number of data points N_{\min} under which further partitions become untrustworthy;
- the distortion measure given by the within-cluster sum of squares of the data vectors associated with the node is below a minimum threshold $WCSS_{\min}$;
- the maximum height of the tree h_{\max} or, equivalently, the maximum binary code length have been reached.

2) *Scalable in-memory clustering*: The proposed indexing mechanism is based on a scalable version of the TSVQ algorithm. The method builds on a distributed implementation of the k -means algorithm to estimate the optimal space partitioning, allowing us to parallelize the clustering operations among the worker nodes of the computing cluster.

A schema of a possible MapReduce implementation of the k -means algorithm is shown in Fig. 7. Input data points are stored on a distributed file system as key/value pairs. The mapping phase consists of the so called *assignment* step. At the i th iteration, the set of k centroid estimates $\mathbf{y}_1, \dots, \mathbf{y}_k$ is globally broadcast to map workers. Each map function has in input a group of data points and iteratively evaluates the squared Euclidean distance between the data points and each of the centroids vectors. The outputs consist of a key/value pair for each data point, with the value being the input vector and the key being its nearest centroid. Each map worker can perform a combine operation, coupling to each centroid key the sum and the number of its associated vectors. This last local combination operation is in principle optional, yet it allows to avoid using expensive network resources by compressing the information generated by each node. The *update* step of the k -means algorithm is finally

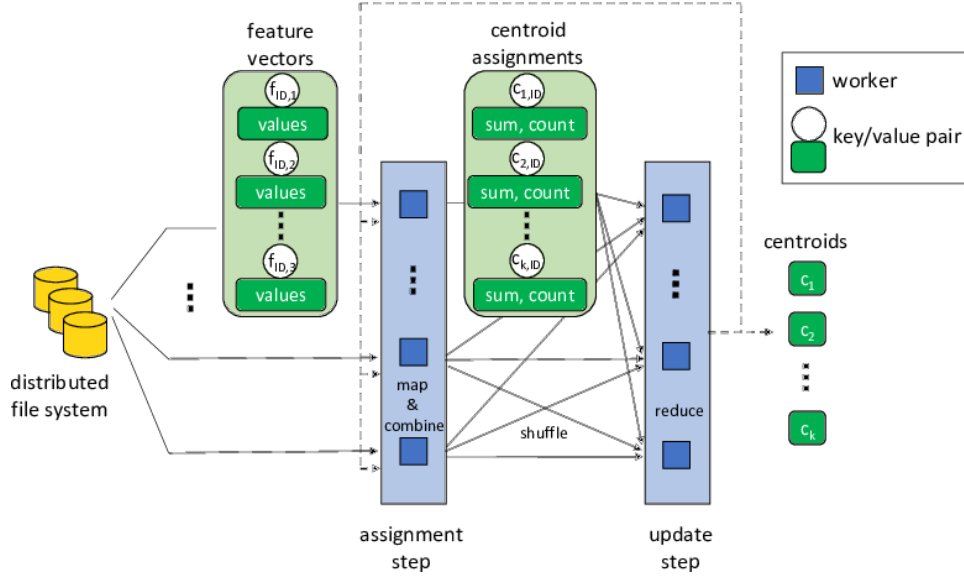


Fig. 7. General MapReduce implementation of the k -means algorithm. The assignment and the update step can be implemented by a sequence of map&combine and reduce phases. Tile descriptors stored in a distributed storage system are given as input to mappers together with the current centroids estimates. Each mapper performs the assignment step, putting out a key/value pair for each input vector, with the value being the input vector and the key being its nearest centroid. A combiner performs then a group-by-key, coupling to each centroid-key the sum and the number of its associated vectors. Reducers finally aggregate the sets of key/value pair to provide updated estimates of centroids.

implemented by the reduce workers. Each reduce function takes in input the grouped-by-key intermediate key/value pairs and computes for each centroid the average of the associated values

$$\mathbf{y}_i^{(t+1)} = \frac{1}{|D_i^{(t)}|} \sum_{\mathbf{x}_j \in D_i^{(t)}} \mathbf{x}_j, \quad (16)$$

where t is the current iteration and D_i is the set of vectors belonging to centroid \mathbf{y}_i . Each average value represents an updated centroid to be used in successive iterations. The algorithm stops when the update variations are below a given threshold.

The assignment step of the algorithm requires to repeatedly access data points and to evaluate the distances between each value vector and the cluster centroids:

$$\begin{aligned} d(\mathbf{x}, \mathbf{y}_i^{(t)}) &= \|\mathbf{x} - \mathbf{y}_i^{(t)}\|_2^2 \\ &= \|\mathbf{x}\|_2^2 - 2\mathbf{x} \cdot \mathbf{y}_i^{(t)} + \|\mathbf{y}_i^{(t)}\|_2^2. \end{aligned} \quad (17)$$

As the data points are fixed during iterations, the first term in the second line of the expression in Eq. 17 does not need to be recomputed at each iteration of the algorithm. Furthermore, also the third term, i.e. the centroid squared Euclidean norm, is repeatedly used inside each single iteration. This indicates the advantage of performing operations in-memory and avoiding, at each iteration, costly operations like both recomputing invariable quantities and storing and reading intermediate results on low performing devices. While standard cluster computing frameworks like Hadoop [19] routinely materialize to the distributed file system intermediate results produced by the directed acyclic graph of operations, with massive usage of cluster resources especially in the case of iterative processing schemes, in-memory ones like Apache Spark [21], [20] avoid this step and have therefore been selected as the base for our implementation.

Further improvements can be obtained accounting for the sparsity of the feature vectors. If the vectors are sparse, the computational complexity of performing k -means decreases from $O(ndki)$ to $O((n_{nz} + n)ki)$, where n is the number of data points, d is the data space dimensionality, k is the number of clusters, i is the number of iterations needed for convergence and n_{nz} is the (average) number of non-zero elements in each data point.

3) *TSVQ Complexity*: The implemented TSVQ for scalable indexing processing works with a fixed number of clusters, $k = 2$. This choice allows the association of binary codes to the leaves of the tree, each with a length equal to the level of the corresponding leaf. As each leaf corresponds to a partition centroid, all the points will be labeled with the binary code of the leaf they belong to. Given the tree structure generated by the TSVQ algorithm and a query vector, retrieving the nearest feature vectors consists of traversing the tree from the root node and choosing, at each branch, the nearest child node, until a leaf is reached. The retrieved data points are the ones with the same codeword as the final leaf.

There are two main motivations behind the choice of the TSVQ algorithm for the index formation. First, this approach leads to the definition of a binary indexed tree, a structure that allows efficient lookup and update operations. In particular, search and modification operations can be executed in constant or $O(\log n)$ logarithmic average times, instead of the $O(n)$ linear

TABLE I
CLUSTER WORKER NODE CHARACTERISTICS.

ID	Processor			vCPU	Memory (Gb)	Network bandwidth (Mbps)
	Model	Frequency (GHz)	Cache size (MB)			
m1.small	E5-2651 v2	1.80	30	1	1.70	~ 300 (Low)
m1.medium	E5-2650	2.00	20	1	3.75	~ 850 (Moderate)
m1.large	E5-2651 v2	1.80	30	2	7.50	~ 850 (Moderate)
m1.xlarge	E5-2651 v2	1.80	30	4	15.00	~ 1000 (High)
c1.medium	E5-2651 v2	1.80	30	2	1.70	~ 850 (Moderate)
c1.xlarge	E5-2650	2.00	20	8	7.00	~ 1000 (High)

access times required by linked lists. Second, as described in Sec IV-B2, the k -means algorithm can be reimplemented to scale to massive sets of data. In addition, scalable implementations of seeding strategies have been recently proposed, allowing for faster runs and more robust searches with respect to suboptimal solutions [1].

C. Search and interactive visualization subsystems

The interactive visualization subsystem provides users with a query interface to the search engine where they can select or upload a query/example image tile and set metadata constraints for the search (geographic coverage, time of acquisition, sensor parameters, mission annotations). The search subsystem is then invoked to transform the user provided examples to a form suitable for query processing, i.e. performing any pre-processing and feature extraction operations on them, and then to determine the best matching tiles based on the content index. Finally, the interactive visualization subsystem receives the identifiers of the best matching tiles and displays the results.

V. EXPERIMENTS

We considered full polarimetric SAR data products from NASA JPL's UAVSAR image archives. In the public archive, six images for each product are made available, that is the cross-products of the 4 Single Look Complex files representing the measurements of the scattering matrix S_{hh} , S_{hv} , S_{vh} and S_{vv} . Each product is available in the ground-projected polarimetric format (equiangular geographic projection, 6-by-6 m pixels resolution). In the experiments, data ingestion has been carried out by instantiating five nodes on the elastic cloud computing infrastructure by using the Amazon EC2 service¹. In the tiling phase, each image is divided in 512-by-512 pixel sized square patches, corresponding to approximately 3-by-3 km areas. Then, for each tile, the feature extraction module performs pixel-by-pixel decomposition procedures, forms multi-dimensional histograms and pushes the vectorized tile descriptors to a durable cloud-based infrastructure, the Amazon S3 service², in which data are redundantly stored across multiple facilities. The reason we use such a durable storage service is that the results of the ingestion need to be reused in multiple experiments. In this work, we performed the ingestion of 43 UAVSAR products, amounting to about 0.67 Terabytes of data, from which we obtained 12950 tiles, with a total size for the descriptors of 2.4 GB. The replication factor in the HDFS filesystem has been set to three in all experiments.

A. Scalability analysis

An experimental analysis of both vertical and horizontal scalability has been conducted by varying the cluster configuration. Horizontal scalability is the ability to face increases in workload by adding worker entities to an application environment. Fig. 8 shows the results of the tests carried out by varying the number of nodes in the cluster. The cluster instances are of type *m1.medium* (see Tab. I). The relation between the number of nodes and the times needed for data processing can be expected to be linear. For this reason a least squares linear fitting has been carried out on the results and the presence of possible outliers has been assessed by Bonferroni method with familywise significance threshold $\alpha = 0.05$. The first plot refers to the time needed to transfer data from the durable (S3) storage to the cluster distributed storage (HDFS). Since both storage systems are distributed, data transfer can be carried out in parallel and the linear relation is quite well satisfied. The second plot refers to the indexing running times. In this case, an outlier is detected for the simplest cluster, with two nodes, while the linear dependence is well satisfied for all larger configurations. The reason for which in the two-nodes configuration the performance is worse and falls outside the linear dependence relation can be attributed to an overall shortage of memory. This causes the loss of already computed in-memory intermediate results, thereby forcing their recomputation from HDFS and degrading the overall system performance.

The advantage of performing in-memory computations during the indexing procedure is better exemplified in Fig. 9. Standard frameworks for data analytics like Apache Hadoop would require a write operation at each iteration in order to save intermediate results on the distributed file system. This operation involves overhead, as data must be distributed and replicated across

¹<http://aws.amazon.com/ec2/>

²<http://aws.amazon.com/s3/>




Fig. 8. Horizontal scalability analysis. Black squares are running times in minutes by varying the number of instances in the cluster, while dashed lines are linear regression tests obtained by discarding possible outliers as determined by Bonferroni test ($\alpha < 0.05$). The cluster instances are of type *m1.medium* (see Tab. I). In (a) the times needed for the data transfer from the durable storage (S3 service) to the cluster distributed file system are shown, while (b) refers to the times needed for the indexing procedure.




Fig. 9. Running times in minutes of the indexing procedure obtained by Apache Hadoop, where performances are bounded by replication, serialization and disk I/O procedures needed for data sharing across parallel jobs, and Apache Spark, that is based on a shared memory-based architecture. In the experiments, the cluster is composed by 3 nodes and the database contains 12950 tiles.

the nodes, thereby increasing the total network traffic and slowing the computation further. The results in Fig. 9 refer to an experiment performed with a cluster of 3 *m1.medium* nodes. In this case, exploiting in-memory computations improves indexing performance by almost 90%.

Vertical scalability is the ability to increase the computation capabilities of the cluster by improving the capacity of existing hardware or software. In theory, the more the resources – as for instance memory, network bandwidth, number of virtual CPUs – we allocate in the cluster, the more the work that can be carried out in a certain amount of time. However, increasing the hardware capacities works just up to a certain point, as the best configuration depends on the algorithm, on the user needs and on their means. In Fig. 10 we show an analysis of the indexing running times by varying the kind of instances allocated in the cluster, as reported in Tab. I. The dimensionality of the cluster is kept fixed to three. A first observation in this experiments is about the cheapest kind of instance, the *m1.small*, which performs considerably worse than all other cases. The reason is here a combination of scarceness of both memory and network resources. On the other hand, the improvements in the other experiments are mainly attributable to the increase in the number of virtual CPUs allocated.

As for the system response time to queries, it is in the order of seconds and essentially independent of variations of the cluster settings.

B. Cost analysis

We also performed a cost-benefits analysis, in order to evaluate whether investing on more resources brings effective advantages, and to tune the system according to this analysis. Results are reported in Fig. 11 and refer to the indexing phase. The two plots show the relation between the improvement percentage of the indexing running time compared to the cost variation percentage obtained by varying the resources allocated in the cloud, with respect to the worst case, as described below. The analysis for a variable number of nodes in the cluster is shown in Fig. 11a, where the instances are of type *m1.medium*




Fig. 10. Vertical scalability analysis. Running times in minutes of the indexing procedure by varying the kind of instances in the cluster (see Tab. I). In the experiments, the cluster is composed by 3 nodes.

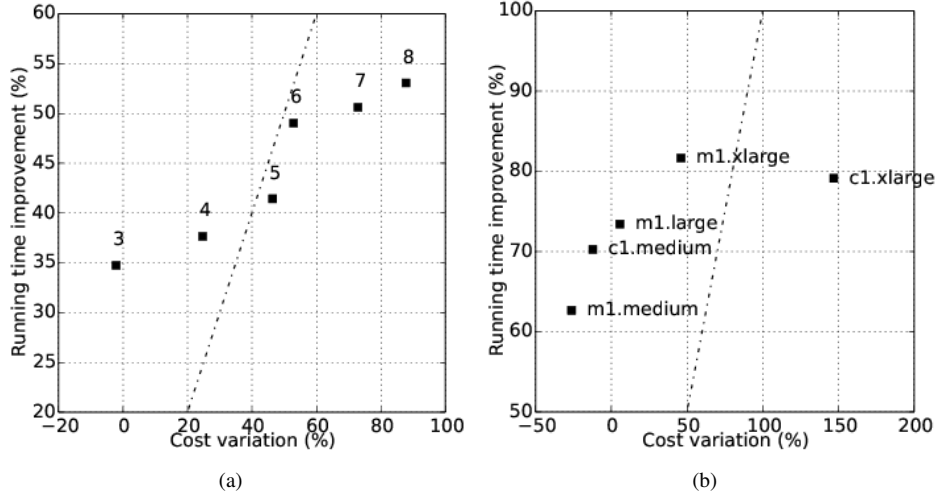


Fig. 11. Cost-benefit analysis. Improvement percentage of the indexing running time versus the cost variation percentage obtained by varying the cluster dimensionality (with *m1.medium* instances), in (a), and the kind of instance used (with three nodes), in (b). The comparison is made with respect to the worst case, i.e. to a cluster composed by two nodes for (a) and to a cluster with *m1.small* instances for (b). The straight dashed line in both figures is a line with unitary slope, representing the locus of the points where time improvements correspond to equal cost increments.

and the worst case for comparison is the use of two nodes (see Fig. 8). This test shows that adding more nodes to the cluster diminishes the indexing running times, but increasing investments are not balanced by equivalent increases of performance. The best trade-off has been obtained for the three nodes configuration for the specific case of the data volumes considered in the experiment.

Similar considerations hold for the analysis shown in Fig. 11b, which refers to the use of different types of instances, with a cluster dimensionality fixed to three and the worse case for comparison is the use of *m1.small* instances (see Fig. 10). In this case, positive balances are obtained for both cases *m1.medium* and *c1.medium* instances, while the worse case happens for *c1.xlarge* instances.

C. Quality assessment

Evaluating the quality of content-based retrieval systems is a challenging task. In the literature, quality assessment is not always present or is based on non-public databases of categorized images. Evaluating the retrieval performance on massive archives of images is even more difficult. A first idea to automatically assign labels to our dataset was to exploit publicly available worldwide land cover data, as those provided by ESA's Globcover project, obtained by MERIS sensor data. UAVSAR products were coregistered with land cover maps and image patches were obtained from both sources. Land cover tile descriptors consisted of the vectorized class labels and a separate index for such dataset was created. We then tried to evaluate retrieval performance by querying the two search engines with coregistered example-query images, i.e. by polarimetric SAR and landcover images respectively. However very poor results were obtained, due to differences in the spectral bands of the sensors (15–30 cm for L-band UAVSAR products and 390–1040 nm for MERIS sensor), since the information collected about a target is different in the two cases, and to the variability in their spatial resolution (3 m and 300 m respectively).

Therefore we resort here to visual inspection analysis. To this aim, the system has been repeatedly queried by different example-based queries. The results of ten experiments, each characterized by different target characteristics, are shown in

Tab. II, where the query tile is shown together with quick looks of the top ten tiles displayed by the interactive visualization subsystem. We consider that the retrieved tiles are highly consistent with the query images in terms of content similarity. Better performances might be achieved by enriching the tiles content description by textural and shape descriptors as well as more advanced descriptions for metric resolution data based on signal decomposition in fractional frequency domains as in [18], [17].

VI. DISCUSSION AND CONCLUSIONS

Content-based EO image retrieval is an open research area that is currently receiving particular attention because of the fast and continuous increase in the volume of acquired product repositories. While several systems for image mining have been proposed, they often lack the capability of managing and accessing massive amounts of data because of their architectural and algorithmic approaches oriented to limited archive volumes.

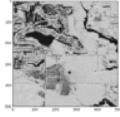



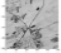



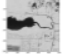







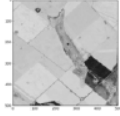















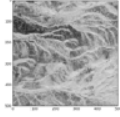




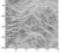
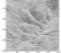

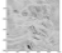

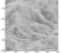





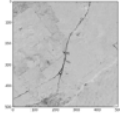


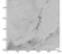

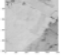
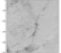









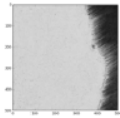


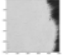


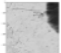


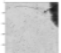






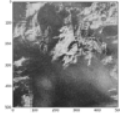







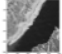




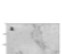


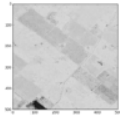

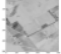



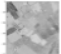

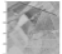







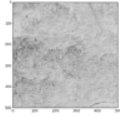
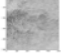
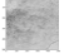













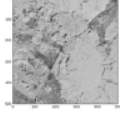





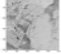









We propose using inexpensive computing clusters instantiated on public Infrastructure-as-a-service cloud platforms in order to ingest and index Petabyte-scale remote sensing image archives with the aim of implementing query-by-example retrieval functionalities on them. To this end, we developed and tested a full prototype operating on a subset of polarimetric SAR products from the NASA JPL UAVSAR archive.

The ingestion component of the system is centered on an elastic load balancer instantiating virtual resources in elastic platforms based on the occupation of input queues, in order to distribute ingestion processes on multiple machines. Furthermore, the indexing component is a parallel/scalable indexing algorithm built upon a cluster-computing framework to enable the prototype to face the difficulties of performing large scale processing on massive image archives. The use of public Infrastructure-as-a-service cloud platforms has proved both flexible, by providing dynamic scaling of resources, and advantageous, as it allows economies of scale. Performances of retrieval measured have proved both the efficiency of the system in terms of response times and its effectiveness in terms of retrieval quality. Although proposed for the analysis of polarimetric SAR image archives, the proposed approach is general and can be expanded to ingest archives from different sensors and missions.

REFERENCES

- [1] Bahman Bahmani, Benjamin Moseley, Andrea Vattani, Ravi Kumar, and Sergei Vassilvitskii. Scalable k-means++. *Proceedings of the VLDB Endowment*, 5(7):622–633, 2012.
- [2] Shane R Cloude and Eric Pottier. A review of target decomposition theorems in radar polarimetry. *Geoscience and Remote Sensing, IEEE Transactions on*, 34(2):498–518, 1996.
- [3] PC Cosman, C Tseng, R Gray, RA Olshen, LE Moses, HC Davidson, CJ Bergin, and EA Riskin. Tree-structured vector quantization of CT chest scans: image quality and diagnostic accuracy. *Medical Imaging, IEEE Transactions on*, 12(4):727–739, 1993.
- [4] Jeffrey Dean and Sanjay Ghemawat. MapReduce: simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.
- [5] Ian Foster. *Designing and building parallel programs*. Addison Wesley Publishing Company, 1995.
- [6] Allen Gersho and Robert M Gray. *Vector quantization and signal compression*. Springer, 1992.
- [7] Nikolas Roman Herbst, Samuel Kounev, and Ralf Reussner. Elasticity in cloud computing: What it is, and what it is not. In *Proceedings of the 10th International Conference on Autonomic Computing (ICAC 13)*, pages 23–27. USENIX, 2011.
- [8] M Klaric, G Scott, Chi-Ren Shyu, C Davis, and K Palaniappan. A framework for geospatial satellite imagery retrieval systems. In *Geoscience and Remote Sensing Symposium, 2006. IGARSS 2006. IEEE International Conference on*, pages 2457–2460. IEEE, 2006.
- [9] Manolis Koubarakis, Michael Sioutis, George Garbis, Manos Karpathiotakis, Kostis Kyzirakos, Charalampos Nikolaou, Konstantina Bereta, Stavros Vassos, Corneliu Octavian Dumitru, Daniela Espinoza-Molina, et al. Building virtual earth observatories using ontologies, linked geospatial data and knowledge discovery algorithms. In *On the Move to Meaningful Internet Systems: OTM 2012*, pages 932–949. Springer, 2012.
- [10] Jong-Sen Lee and Eric Pottier. *Polarimetric radar imaging: from basics to applications*. CRC press, 2009.
- [11] Michael McCool, James Reinders, and Arch Robison. *Structured parallel programming: patterns for efficient computation*. Elsevier, 2012.
- [12] Diana Moise, Denis Shestakov, Gylfi Gudmundsson, and Laurent Amsaleg. Terabyte-scale image similarity search: experience and best practice. In *Big Data, 2013 IEEE International Conference on*, pages 674–682. IEEE, 2013.
- [13] Marco Quartulli and Igor G Olaizola. A review of EO image information mining. *ISPRS Journal of Photogrammetry and Remote Sensing*, 75:11–28, 2013.
- [14] Paul A Rosen, Scott Hensley, Kevin Wheeler, Greg Sadowy, Tim Miller, Scott Shaffer, Ron Muellerschoen, Cathleen Jones, Howard Zebker, and Soren Madsen. Uavsar: a new nasa airborne sar system for science and technology research. In *Radar, 2006 IEEE Conference on*, pages 8–pp. IEEE, 2006.
- [15] Konstantin Shvachko, Hairong Kuang, Sanjay Radia, and Robert Chansler. The Hadoop distributed file system. In *Mass Storage Systems and Technologies (MSST), 2010 IEEE 26th Symposium on*, pages 1–10. IEEE, 2010.
- [16] Chi-Ren Shyu, Matt Klaric, Grant J Scott, Adrian S Barb, Curt H Davis, and Kannappan Palaniappan. GeoIRIS: Geospatial information retrieval and indexing system-content mining, semantics modeling and complex queries. *Geoscience and Remote Sensing, IEEE Transactions on*, 45(4):839–852, 2007.
- [17] Jagmal Singh and Mihai Datcu. SAR image categorization with log cumulants of the fractional Fourier transform coefficients. *IEEE transactions on geoscience and remote sensing*, 51(12):5273–5282, 2013.
- [18] Marc Walessa and Mihai Datcu. Model-based despeckling and information extraction from SAR images. *Geoscience and Remote Sensing, IEEE Transactions on*, 38(5):2258–2269, 2000.
- [19] Tom White. *Hadoop: the definitive guide: the definitive guide*. " O'Reilly Media, Inc.", 2009.
- [20] Matei Zaharia, Mosharaf Chowdhury, Tathagata Das, Ankur Dave, Justin Ma, Murphy McCauley, Michael J Franklin, Scott Shenker, and Ion Stoica. Resilient distributed datasets: a fault-tolerant abstraction for in-memory cluster computing. In *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*, pages 2–2. USENIX Association, 2012.
- [21] Matei Zaharia, Mosharaf Chowdhury, Michael J Franklin, Scott Shenker, and Ion Stoica. Spark: cluster computing with working sets. In *Proceedings of the 2nd USENIX conference on Hot topics in cloud computing*, pages 10–10, 2010.

TABLE II
CONTENT-BASED RETRIEVAL RESULTS

ID	Query image	Retrieved tile thumbnails (first 10)									
Q1											
											
Q2											
											
Q3											
											
Q4											
											
Q5											
											
Q6											
											
Q7											
											
Q8											
											
Q9											
											
Q10	